arXiv:1607.02399v1 [cs.DB] 8 Jul 2016

Translating Bayesian Networks into Entity Relationship Models
(Extended Version)*

Frank Rosner
Global Data and Analytics
Allianz SE
Germany

Alexander Hinneburg
Computer Science
Martin-Luther-University
Halle-Wittenberg
Germany

Abstract

Big data analytics applications drive the convergence of data management and machine learning. But there is no conceptual language available that is spoken in both worlds. The main contribution of the paper is a method to translate Bayesian networks, a main conceptual language for probabilistic graphical models, into usable entity relationship models. The transformed representation of a Bayesian network leaves out mathematical details about probabilistic relationships but unfolds all information relevant for data management tasks. As a real world example, we present the TopicExplorer system that uses Bayesian topic models as a core component in an interactive, database-supported web application. Last, we sketch a conceptual framework that eases machine learning specific development tasks while building big data analytics applications.

1 Introduction

The implementation of a big data analytics application requires to join data management software with machine learning tools. However, the fields of data management and machine learning developed quite different models and notations. The former frequently uses entity-relationship models (ERM) while the latter uses probabilistic graphical models to communicate key concepts. In this paper, we pick Bayesian networks (BN) as a widely used graphical notation for machine learning models. Note that the presented ideas can be transferred to

*This document is an extended version of a paper published in the Proceedings of the 35th International Conference on Conceptual Modeling, ER 2016. In addition to a more detailed discussion of the translation method presented in the conference version, this extended version provides a description of a case study that applies the method as well as first ideas of a conceptual framework for developing big data analytics applications.

other common graphical notions like undirected Markov models or factor graphs as well.

The notations of ERMs and BNs are designed to serve the needs of the respective fields. They are stressing information relevant in the particular domain while visually suppressing less important details. E.g. an ERM highlights the existence and cardinalities of relationships between distinguishable entities. On the other hand, a BN represents a joint probability distribution as a factorization of several hierarchically linked conditional probabilities. Even while both kinds of graphical notations show many details of the data, information explicit on one side remains implicit on the other one and vice versa — there is no natural understanding of the two worlds. However, a common conceptual description of the contribution from both worlds is crucial for the successes of big data analytics projects.

The data management part of a big data analytics project typically represents more details of the data than the machine learning part. Therefore, it is reasonable to translate the machine learning part to the conceptual language of data management. Attempts into this direction include machine learning libraries with APIs in one or multiple programming languages [17, 9, 16], and new declarative languages or extension of existings ones [2, 15, 7, 1, 12]. However, none of these approaches solves the problem of integrating the information about machine learning that is relevant for data management into the conceptual view of this side. The advantages of a formal conceptual view of machine learning models integrated into the conceptual view of the data management side would be (i) no black box behind an abstract API in the data management model and (ii) developers from the data management side understand the basic in- and outputs of the machine learning part.

We propose a rule-based method to translate a graphical BN model in plate notation into an entity relationship model. Such ERM can be easily integrated into the overall ERM of the whole application. As an example, we look at topic modeling of documents [3]. Latent Dirichlet allocation (LDA) [4] – shown in Figure 9a – is one of the most popular topic models. Although not limited to this application, it is often used to find a hidden structure in text documents, called *topics*. Topics are modeled as probability distributions over a vocabulary. They are often presented as word lists each ordered by descending probabilities.

The given data in this example include documents consisting of word tokens. Each token corresponds to a particular occurrence of a word from the vocabulary. Documents, tokens and words have additional information attached like title, publication date or part-of-speech (POS) tags. The entities and relationships about the given data are shown in Figure 1 (left). The BN describing the LDA topic model shown in Figure 9a represents the given word tokens $\vec{d}_{nm}$ as shaded circles, which indicates random variables with given values. The hidden random variables represented by empty circles are the token-topic assignments $\vec{z}_{nm}$, document specific topic proportions $\vec{\theta}_n$ and topic-vocabulary distributions $\vec{\mu}_k$. For those variables, either expectations or probable value assignments are computed during machine learning inference. The black dots represent fixed

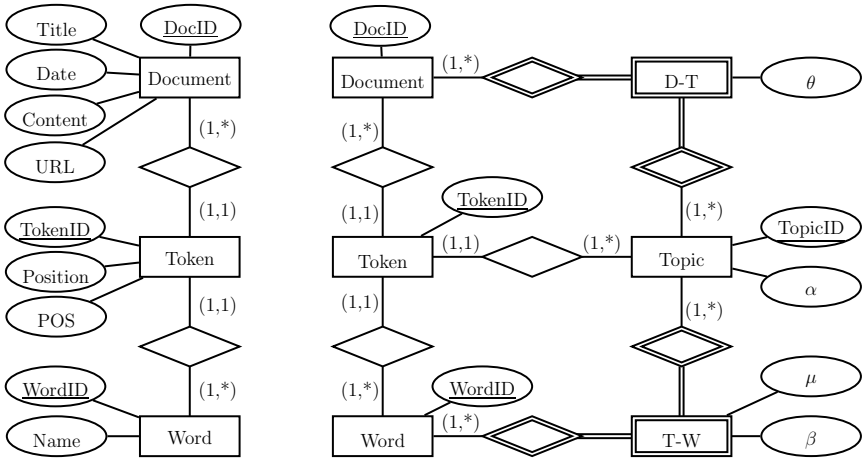

Figure 1: ERM of given Data (left) and translated ERM for LDA (right).

hyper-parameters that determine the prior distributions for the hidden variables. The descriptions of the boxes (plates) correspond to the entities `Document` and `Token`. However, some entities like `Word` do not appear in the BN as they are implicit in the definition of vectors. Further, the plate named `Topics` introduces a new entity.

The result of the whole translation is shown in Figure 1 (right). The ERM of the BN represents all relevant entities with respective primary keys. Further, the relationships between topics on one side and documents, tokens and vocabulary words on the other side are shown with their respective cardinalities. The ERMs about the given data and the translated one from BN can be combined into a single one by merging the matching entities. We believe that such an overall ERM helps to improve the efficiency of the development process as now the developers on the data management side can see what data is needed and contributed from the machine learning part with respect to relational aspects.

Our main contribution is a rule-based translation from BN in plate notation to ERMs, described in Section 2. Our method provides semantic guidelines for building a conceptual representation of a BN that addresses the needs of the data management side of a big data analytics project. However, a BN is not a unique way to describe a probabilistic model, i.e. the same probabilistic model can be described by multiple BNs that differ in complexity. Therefore, our proposed translation constructs an intermediate atomic plate model (APM) in several steps. It gradually uncovers implicit information not represented in the original BN. Further, the subsequent translation from an APM to an ERM can include different probabilistic relationships between the generated entities.

We demonstrate the method in the real world example of the TopicExplorer in Section 3. Based on this, we describe our vision of a conceptual framework that uses pre-translated BNs as a library of ERM snippets in Section 4. Such library could be used by data management developers to conceptually include

machine learning methods into analytics applications. We believe that software development of big data analytics applications could benefit from machine learning implementations that are attached to the pre-translated BNs. Last, we discuss related work in Section 5 and conclude the paper in Section 6.

2 Translation of Bayesian Network in Plate notation

BNs mainly describe data with random variables at the level of data items, e.g. word tokens. The BN for a topic model in Figure 9a has an observed random variable $\vec{d}_{nm}$ for each of the M_n tokens of the nth document. As plotting all data items one by one is not possible, the visual notation of plates is used. A plate groups a set of random variables sharing an index set. Due to this, plates convey information about entities and relationships. However, some random variables are implicitly denoted in BNs: (i) multidimensional vector notation of random variables and (ii) functions that implicitly describe data transformation coupled with joins. Therefore, we propose a stepwise approach to transform a given BN in plate notation into a well-formed ERM. This is done in three steps:

1. Make implicit relational information explicit: The resulting model is called an atomic plate model (APM).
2. Convert the APM into an ERM based on graphical rules.
3. Reduce the ERM to avoid translation artifacts.

We use the standard ERM notation [6] with min-max cardinalities [8, p. 82]. We call an ERM *well-formed* iff it (1) is syntactically correct, (2) is explicit and (3) does not have redundant constructs. Explicitness means that all real world constructs which have a corresponding construct in the ERM notation are modeled using those. Thus, a well-formed ERM does not contain explicit foreign key attributes but uses relationships instead. Having no redundant constructs means that no entity or relationship is duplicate, i.e. semantically expressing the same thing. However, those duplicates may appear as intermediate results of the translation procedure. The following subsections offer detailed explanations of all steps to translate a BN to an ERM.

2.1 Construction of Atomic Plate Models

Plate models (BN in plate notation) may contain multidimensional random variables (e.g. vectors or matrices), hidden deterministic data transformations, and relationships. A plate model is converted to an APM by explicitly including those hidden transformations and relationships and any variables associated with it, as well as splitting multidimensional variables into their components.

Figures 2a and 2b show the conversion of plate models containing a vector and a matrix variable, respectively. The original multidimensional variable is

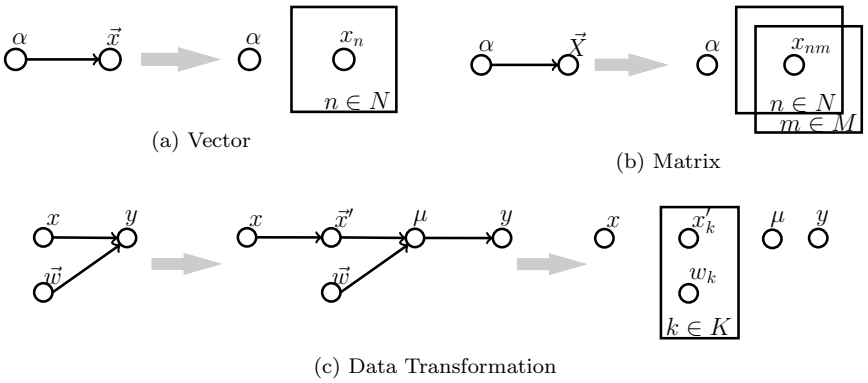

(a) Vector

(b) Matrix

(c) Data Transformation

Figure 2: Conversion of Bayesian networks to atomic plate models

represented in the APM by a plate containing an indexed component vertex. The multidimensional variables, e.g. $\vec{x} \in \mathbb{R}^{|N|}$ and $\vec{X} \in \mathbb{R}^{|N| \times |M|}$, are converted into individual component variables x_n and x_{nm}, surrounded by plates of the corresponding dimensions that are described by the respective index sets N and M.

In this translation, edges are discarded. If they would be preserved, they conveyed wrong semantics about the conditional probabilities of the BN after the decomposition. However, it is not necessary to preserve this information in the conversion process to a database model, as it will not be used as a probabilistic model anymore. By discarding all edges after the conversion process we avoid confusion as they might not represent conditional probability distributions in the APM.

Detecting the deterministic relationships and data transformations hidden in BNs is a bit more subtle. Figure 2c illustrates this on the example of polynomial regression. The one-dimensional input x is transformed into a vector $\vec{x}'$ of several inputs $x'_k = x^k$ by taking different powers $k \in K \subset \mathbb{N}$. All powers are multiplied with weights that are components of the vector $\vec{w} \in \mathbb{R}^{|K|}$. Finally, the weighted powers are summed up and this sum is used as the mean μ of a normal distribution that governs a univariate random variable y. As we introduced another multidimensional variable $\vec{x}' \in \mathbb{R}^{|K|}$, splitting all multidimensional variables yields a common plate with index set K including all x'_k and w_k.

2.2 Translation of APM to ERM

After converting a plate model to an APM, it is translated to an ERM. For this step, we adapt the mapping from plate models to DAPER models [11]. We defer the discussion of the differences between our mapping and DAPER models to the related work section. First, we state the basic translation rules. Then, we consider some special cases of model constraints and show the influence on the

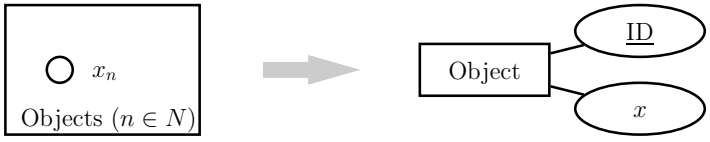

Figure 3: Translation of a plate containing a variable x_n to an entity type having an artificial key ID and an attribute x.

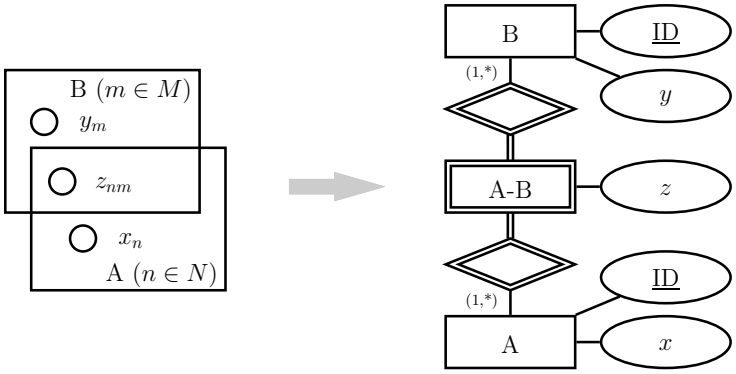

Figure 4: Translation of a plate intersection to an association entity of a many-to-many relationship. The variable z_{nm} residing in the intersection is represented as the attribute z of the association entity.

resulting ERM.

Translate plates to entity types. Each plate is represented as an entity type. Usually there is an index set associated with each plate. Each entity type gets an artificial key (ID) that enumerates the index set. Figure 3 illustrates the translation of a plate, representing a set of Objects N, to the Object entity type.

Translate plate intersections to relationships. In general, plate intersections represent many-to-many relationships between the corresponding entity types. In contrast to [11], we express all relationships as association entity types [8, pp. 86-88]. This allows a generic translation procedure that can easily be used for n-ary relationships. Figure 4 shows the translation of two entity types having a binary relationship.

Translate variables to attributes. The translation of attributes depends on the number of plates surrounding them. If a variable is surrounded by exactly one plate, the entity type of that plate gets an additional attribute representing this variable (see Figure 3). If a variable resides inside multiple plates, it becomes an attribute of the corresponding association entity (see Figure 4). Variables that are associated with no plate are assigned to an artificial entity type called Global (see Figure 5). There exists only one entity of type Global.

Figure 5: Translation of a single variable that is not located in any plate. It will be associated with the artificial entity type Global that has only one entity.

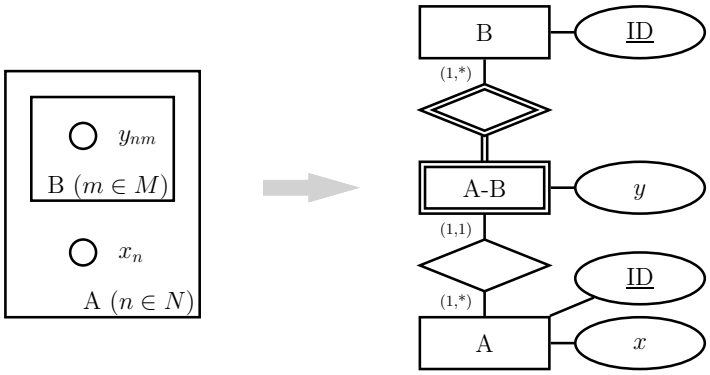

Figure 6: Translation of nested plates to an association entity. The entity type of the covering plate will have no weak relationship to the association entity. This reflects that every instance of B does belong to exactly one instance of A.

Translate nested plates to one-to-many relationships. If a plate is nested in another plate, the resulting relationship has one-to-many cardinality instead of many-to-many. This works well in the simple case of a binary relation, but falls short when one or both plates are additionally intersected or nested with further plates. Therefore, we adjust the mapping rule to the concept of association entities as demonstrated in Figure 6, which works in the general case. In case of only binary relationships like in Figure 6, the weak entity is not necessary. Such cases will be fixed in the reduction step that follows the translation to ERM.

Further, we propose two additional transformation rules that consider the effects of constraints for random variables on the resulting ERM and cope with self relationships as a result of matrix or tensor variable translations.

Adjust cardinalities depending on variable constraints. In some applications there are variables that solely express associations between objects. Those variables can be coded by the 1-out-of-K scheme. This basically says that e.g. the vector $\vec{z} \in \{0,1\}^K$ contains only bits of zero and one and further obeys the additional summation constraint $\sum_{k=1}^{K} z_k = 1$, which says that exactly a single bit in the vector is set to one. In case such variable z_k is in the intersection of two plates and there is such a summation constraint over either index set, then the plate intersection should not be translated as a many-to-many rela-

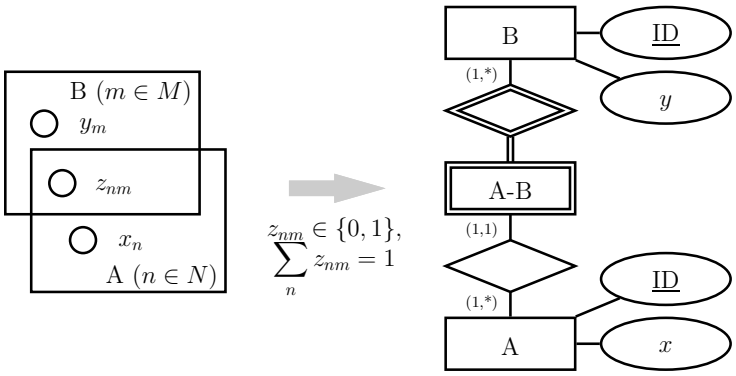

Figure 7: Translation of a plate intersection to an association entity given a 1-of-K relationship variable constraint (z_{nm}). The constraint ensures that every instance of B can only be connected to one instance of A. Instead of representing this variable directly, it will affect the cardinality of the resulting relationship. The entity type whose index set is summed over will not have a weak relationship.

tionship but as one-to-many relationship. However, in this intermediate step of the translation, the one-to-many relationship is expressed by an association entity as in the original many-to-many relationship from the plate intersection to allow the translation to continue with further plate intersections covering z_k. In order to express the one-to-many constraint, the relationship between the entity whose index set is summed over and the association entity is translated as a normal relationship instead of a weak entity relationship. Finally, the constrained attribute z_k of the association entity is removed since now the association information is expressed by the relationships as shown in Figure 7. In case of only two overlapping plates the more complex expression of the one-to-many relationship is simplified in the final reduction step.

Translate overlapping plates with same index set to self relationships. When converting plate models with matrix or tensor variables to APMs, the translation procedure will produce overlapping plates as illustrated in Figure 2b. Given a matrix with equal dimensions $N = M$, this will result in two overlapping plates of the same index set. These are simply translated as a many-to-many self relationship. The matrix components are then represented as an attribute of the resulting association entity.

With this set of rules it is possible to translate any given plate model into an ERM. However, in some cases the resulting model might not be well-formed. It might express simple one-to-many relationships in a complicated way. Therefore, a reduction step is performed as described in the next section.

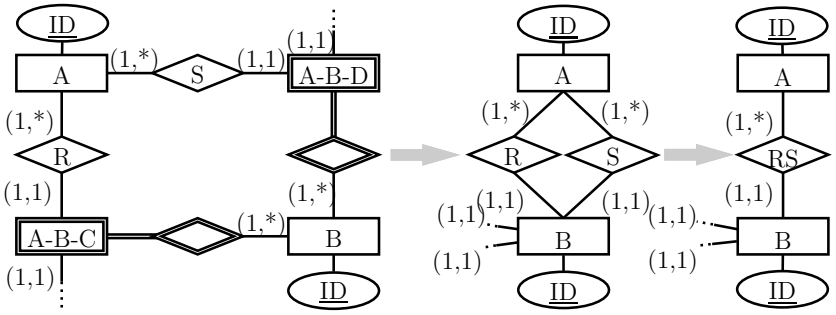

Figure 8: Merging two reduced one-to-many relationships. After A-B-C and A-B-C are merged with B, there will be two relationships between A and B. We merge them if they are equivalent on a semantical level.

2.3 Reduction of Translated Entity-Relationship Models

In order to produce well-formed ERMs from plate models, it is necessary to apply a reduction step after the conversion from APM to ERM. This reduction will make sure that all translation artifacts are eliminated that are caused by straight application of the translation rules. Those artifacts are weak entities without primary key extension and duplicate relationships.

Weak entities without primary key extension may occur in cases like in Figure 6 and 7. Plate intersections are translated to association entities having a weak relationship to all entities that form the intersection. However, 1-of-K coded variables and nested plates turn weak entity relationships into normal ones. In case only one weak relationship is left after the completed translation, the construct is not well-formed when it does not extend the primary key. The weak entity relationship is then a degenerated one-to-one relationship and the weak entity is merged with parent entity.

After performing all possible merges as described above there may be duplicate relationships left. Consider the case in Figure 8. Two weak entities are merged with A and B respectively. Then, the relationships R and S may become redundant. However, merging those is only permitted if the relationships are semantically expressing the same facts. Thus, the translating person needs to evaluate the merging conditions depending on the concrete model. There will be a practical example of this situation in the case study in Section 3.

3 Case Study: Text Topic Modeling

As a practical use case of our approach, we present TopicExplorer, an application to explore document collections using probabilistic topic models [13, 14]. In the first subsection, we explain LDA in detail, which is the probabilistic model behind TopicExplorer and demonstrate the translation of this plate model to an ERM. In the last part, we explain the use case and show typical analyses

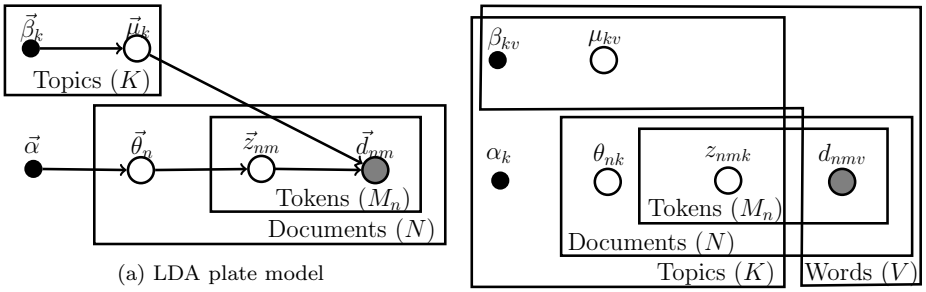

(a) LDA plate model

(b) Atomic LDA plate model

Figure 9: Transformation of the LDA plate model to an APM.

supported by the translated ERM.

3.1 Translation of Latent Dirichlet Allocation Topic Model

LDA models a collection of documents that is indexed by the set N. Each document $n \in N$ consists of a set of tokens M_n that represents the words occurring in this document. The document specific token index sets M_n partition the total index set of tokens M. A token $m \in M_n$ correspondes to exactly one word type v from a vocabulary V. In the Bayesian network, Figure 9a, this information is coded as a bit vector $\vec{d}_{nm} \in \{0,1\}^{|V|}$ that has exactly a single 1 at the index associated with the respective word $v \in V$. Each token $m \in M_n$ is also assigned to a topic $k \in K$. This assignment is coded by the bit vector $\vec{z}_{nm} \in \{0,1\}^{|K|}$, which has a single 1 at the respective topic index. The constraints of the single 1 in both bit vectors are ensured by requiring that the sums of the respective vector elements are one. Furthermore, each topic has its own word distribution parameterized by a vector of positive real number $\vec{\mu}_k \in \mathbb{R}^{|V|}$. The topic proportions per document are represented by a similar vector $\vec{\theta}_n \in \mathbb{R}^{|K|}$. Both vector also sum to one, however, in contrast to the bit vectors, these constraints have no special impact during the translation of the Bayesian network as they are vectors of real numbers. The hyper-parameter vector $\vec{\alpha} \in \mathbb{R}^{|K|}$ and $\vec{\beta}_k \in \mathbb{R}^{|V|}$ regulate the prior distributions for the respecitve hidden parameters $\vec{\theta}_n$ and $\vec{\mu}_k$.

To obtain an ERM for LDA we follow the steps described in Section 2. First, the LDA plate model is transformed into an APM, secondly, this APM is translated to an ERM and last, the ERM is reduced to the final outcome. The transformation to an APM is illustrated in Figure 9. All vector variables are explicitly modeled as repeated variables for their components. The `topic` plate now covers the components of the $|K|$-dimensional vector variables $\vec{z}_{nm}, \vec{\theta}_n$ and $\vec{\alpha}$. A new plate is created for the dimensionality $|V|$ to model the components of $\vec{d}_{nm}, \vec{\mu}_k$ and $\vec{\beta}_k$, representing the vocabulary of the documents.

In the next step, the APM is translated to an intermediate ERM. Figure 10 shows the result. The entities `Document`, `Word`, `Token`, and `Topic` correspond to the respective plates. The association entities `D-T-W`, `D-T-T`, `D-T`, and `T-W`

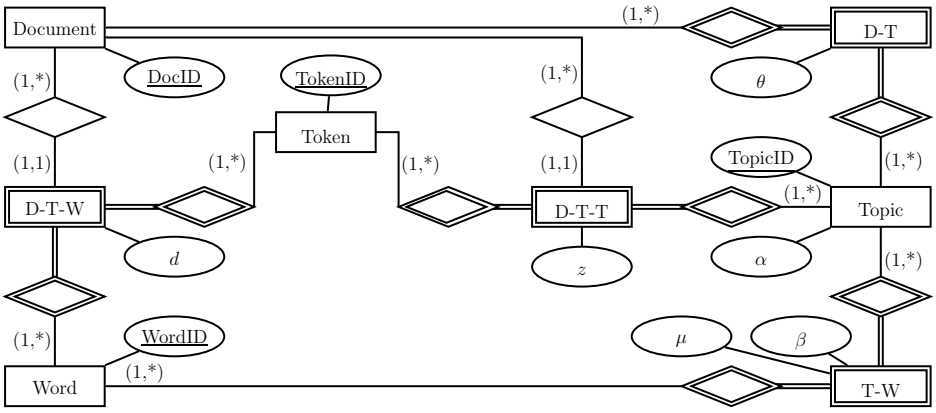

Figure 10: First intermediate ERM for LDA translated from an APM. The entities topic and word still yields non-weak relationships to D-T-T and D-T-W respectively.

represent the plate intersections. The `Token` plate is nested inside the `Document` plate, yielding non-weak relationships from `Document` to `D-T-W` and `D-T-T`. Furthermore, the summation constraints on the bit vector variables $\vec{d}_{nm}$ and $\vec{z}_{nm}$ are considered in the intermediate ERM in Figure 11. This yields a transformation of the relationships between `Word` and `D-T-W`, and `Topic` and `D-T-T`, respectively. These relationships are no longer weak and do not contribute to the primary key of the corresponding association entity. Finally, the intermediate ERM is reduced to the result as shown in the introduction, Figure 1(right). In particular, the `Token` entity type is merged with `D-T-W` and `D-T-T`. The relationships linking to these to entities are then connected to `Token`. Thus, there are now two relationship between Document and Token. As these relationships have the same semantics, they are merged into a single one.

The reduced, final ERM is exactly as one would expect it when designing a data model for LDA from scratch. A document consists of one or more tokens which are of exactly one word type. Each token is assigned to a topic, while one topic can have multiple tokens assigned. The inferred topic mixture for each document is stored in `D-T.`θ, while `T-W.`μ holds the word probabilities for each topic. The hyper-parameter α of the prior for the topic mixture resides as an attribute of the Topic entity type. The parameters for the individual priors for the word distributions are stored in `T-W.`β.

3.2 TopicExplorer

TopicExplorer is a web application that helps users from the humanities to work with topic models, e.g. in a collaboration with the institute for Japanese studies at Martin-Luther-University, we analyzed blog posts about the Fukushima disaster. After crawling relevant blogs, each blog entry is preprocessed by computer

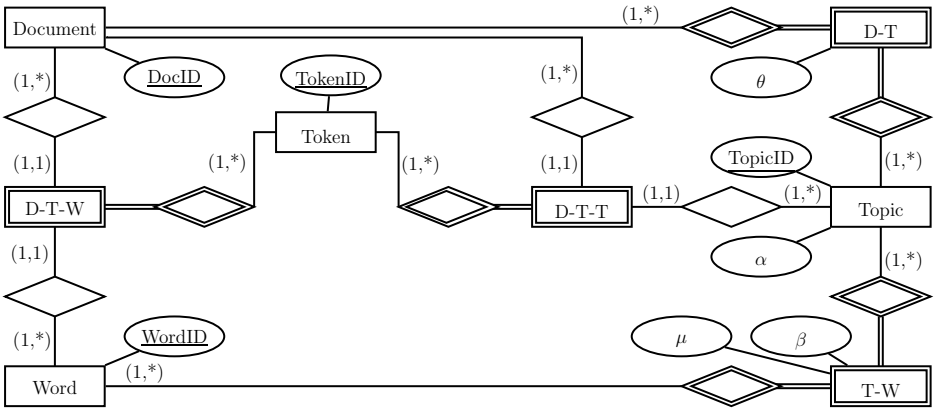

Figure 11: Second intermediate ERM for LDA after adjusting cardinalities and relationships according to the given variable constraints.

linguistic software, which is MeCab (`https://sourceforge.net/projects/mecab/`) in this case, to extracts tokens from full text. The tokens are stored in their lemmatized form together with their part-of-speech tags (e.g. noun, verb or adjective) and their positions in the text. The full text of each document is stored in raw format as well to be able to view it in the user interface. Additionally, the blog post title, URL, and publication date are stored.

TopicExplorer allows user to analyze the topic structure of the documents using interactive visualizations. The visualizations require analysis steps on the given data about the documents join together with results from the topic model. The integrated ERM, see Figure 1, gave the application developer a good idea how to access those data, without needing to understand the machine learning details of a topic model. We present how to derive a few visulizations that are part of the current version of TopicExplorer [13].

Document topic mixture. LDA assigns a vector of topic probabilities stored in `D-T.`θ for each document, called a topic mixture. This could either be presented as a list of topics with decreasing order of probabilities or as a pie chart like in the Topic Model Visualization Engine [5].

Topic Documents. Reversing the idea behind the document topic mixture, one can visualize a topic as a list of the most representative documents for this topic. This is done by joining `Document`, `Token` and `Topic`, grouping by both IDs of documents and topics and counting the number of tokens in each group. For each topic the entries are sorted with decreasing token count, yielding a list of representative documents.

Topic words. As stated above, a topic can be represented as a list of words sorted in descreasing probability (`T-W.`μ). This list can be cut off, yielding a top words visualization for each topic.

Topic frames. Another visualization of topics uses the concept of frames. A *topic frame* consists of a noun and a verb that are assigned to the same topic

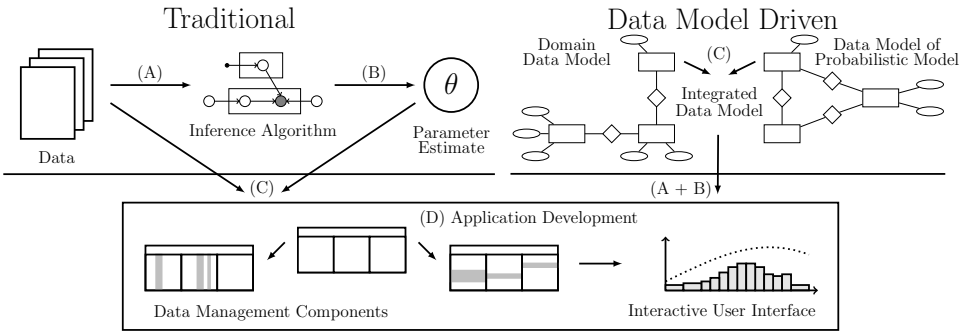

Figure 12: Comparison of traditional development versus data model driven development of big data analysis applications.

and appear close together in the same documents. Thus, instead of visualizing a topic as top words, top frames can be used. Topic frames can be computed using `Token`, `Word` and `Topic`, grouping by topic ID and word IDs of the frame tokens, and counting the number of frames having equal word types per topic.
Topic time. To analyze how the discussion about Fukushima develops, one can visualize how the topics change over time. It is possible to analyze the dynamics of the topics in a postprocessing step after using LDA. To achieve this, entries are grouped depending on their publication date, e.g. monthly. Given those two attributes `Document.Year` and `Document.Month`[1], the `Document` table joined with `Topic` is grouped by the `TopicID`, `Document.Year` and `Document.Month`. Finally, counting the tokens assigned to each topic at each point of time allows plotting a time line for each topic.

4 Conceptual Modeling Framework

Based on the new method for translating probabilistic models to ERMs, we propose a first idea for a new data model driven development approach dedicated to big data analytics applications. Figure 12 visualizes the traditional and our proposed data model driven development approaches. Both address four different tasks, namely (A) gather the data sources and make them available to a probabilistic model, (B) run machine learning components, (C) integrate the data sources with the machine learning output and (D) build the application consisting of data management components and an interactive user interface.

The traditional approach addresses the tasks mainly in sequential order. The first three steps implement data mining process following the CRISP model[10], while the last step addresses standard application development.

The translation method from BN to ERM allows an alternative, data model driven approach. We assume that for a wide spectrum of machine learning

[1]Depending on the implementation of the database, those attributes can be extracted on the fly out of the stored timestamp.

problems abstract, readily developed BNs already do exist. Those BNs could be pre-translated to ERMs to build a library. Thus, conceptual information about the machine learning component is already available when integrating the data sources, task (C). The BN could be treated as just another data source. The entities corresponding to observed variables in the BN, including their respective relationships, have to be matched with those from other available data sources. The matching conceptually defines the interface between data sources and machine learning, task (A). Furthermore, translating the integrated ERM into a (relational) model for a big data framework, e.g. Spark [2], conceptually defines the API between the output of machine learning and the rest of the application, task (B). Depending on the framework, the tasks (A) and (B) could be supported by generation of efficient code for interfaces between data management and machine learning implementations.

As a consequence, the application developers just need knowledge about in- and outputs, and the relationships among the variables in the Bayesian model, but not about probabilistic distributions and dependencies. Thus, our new data model driven approach eliminates unnecessary complexity caused by a lack of compatible conceptual languages on both sides of machine learning and data management. Thereby, it makes the collaboration between machine learning experts and the application developers more direct and offers potential for efficiency optimization.

5 Related Work

Directed acyclic probabilistic entity-relationship (DAPER) models [11] are closely related to our work. However, they are designed to unify probabilistic relational models and plate models. While we discard all facts about probablistic relationships in our translation, DAPER models can still be used to express such relationships. As a consequence, vector variables and other implicit information is not resolved in DAPER models like in the proposed translation to APMs. Thus, DAPER models are not used and also not intended to work as ERMs for conceptional design of data management.

There are several approaches that combine data management with machine learning, however, none of them reaches a comparable conceptual level like ERMs. Hazy [15] provides programming, infrastructure and statistical processing abstractions, the latter are based on Markov logic [7]. This requires a deeper understanding of the machine learning algorithms in order to combine them effectively with data management. Several approaches combine machine learning APIs with SQL [1, 12, 2] or with their own declarative language [17].

Last, data management is combined with machine learning at the level of user interfaces. Examples are Weka [9] and scikit-learn [16], which enable users to quickly select data sources and try different algorithms. Both do not offer an easy way to integrate machine learning results with domain specific meta data. Our approach also contrasts with statistical programming languages and software like R and SAS that just offer programming APIs to data sources and

machine learning algorithms.

Our work is closely related in spirit to a recently proposed conceptual modeling framework work for network analytics [18]. In contrast, we target BNs, but we believe that our vision of a conceptual framework for probabilistic models can yield similar benefits.

6 Conclusion

Our proposed translation shows that modeling an ERM for a given BN is a non-trival task. Knowledge of the translation procedure helps data architects to pose the right questions for machine learning experts to uncover implicit information in a BN. The subsequently proposed framework gives guidelines how to effectively build an integrated conceptual model that includes details about domain specific aspects as well as the machine learning side of a big data analytics application. Future work includes the implementation of the framework and optimizing efficiency when translating integrated conceptual models to a particular implementation.

References

[1] Mert Akdere, Ugur Cetintemel, Matteo Riondato, Eli Upfal, and Stanley B Zdonik. The case for predictive database systems: Opportunities and challenges. In *CIDR*, pages 167–174, 2011.

[2] Michael Armbrust, Reynold S. Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K. Bradley, Xiangrui Meng, Tomer Kaftan, Michael J. Franklin, Ali Ghodsi, and Matei Zaharia. Spark sql: Relational data processing in spark. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, SIGMOD '15, pages 1383–1394, New York, NY, USA, 2015. ACM.

[3] David M. Blei. Probabilistic topic models. *Commun. ACM*, 55(4):77–84, April 2012.

[4] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3:993–1022, 2003.

[5] Allison June-Barlow Chaney and David M Blei. Visualizing topic models. In *ICWSM*, 2012.

[6] Peter Pin-Shan Chen. The entity-relationship model—toward a unified view of data. *ACM Transactions on Database Systems (TODS)*, 1(1):9–36, 1976.

[7] Pedro Domingos and Matthew Richardson. Markov logic: A unifying framework for statistical relational learning. *Introduction to statistical relational learning*, pages 339–371, 2007.

[8] Ramez Elmasri and Shamkant B Navathe. *Fundamentals of Database Systems*. Pearson, 2007.

[9] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. The weka data mining software: an update. *ACM SIGKDD explorations newsletter*, 11(1):10–18, 2009.

[10] Jiawei Han, Micheline Kamber, and Jian Pei. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2011.

[11] David Heckerman, Chris Meek, and Daphne Koller. Probabilistic entity-relationship models, prms, and plate models. *Introduction to statistical relational learning*, pages 201–238, 2007.

[12] Joseph M Hellerstein, Christoper Ré, Florian Schoppmann, Daisy Zhe Wang, Eugene Fratkin, Aleksander Gorajek, Kee Siong Ng, Caleb Welton, Xixuan Feng, Kun Li, et al. The madlib analytics library: or mad skills, the sql. *Proceedings of the VLDB Endowment*, 5(12):1700–1711, 2012.

[13] Alexander Hinneburg, Christian Oberländer, Frank Rosner, and Stefan Pessler. Exploring document collections with topic frames. Submitted to CIKM 2014 Shanghai, China, 2014.

[14] Alexander Hinneburg, Rico Preiss, and René Schröder. Topicexplorer: Exploring document collections with topic models. In *Machine Learning and Knowledge Discovery in Databases*, pages 838–841. Springer, 2012.

[15] Arun Kumar, Feng Niu, and Christopher Ré. Hazy: Making it easier to build and maintain big-data analytics. *Communications of the ACM*, 56(3):40–49, 2013.

[16] Scikit. scikit-learn, 2014. Machine Learning in Python.

[17] Evan R Sparks, Ameet Talwalkar, Virginia Smith, Jey Kottalam, Xinghao Pan, Joseph Gonzalez, Michael J Franklin, Michael I Jordan, and Tim Kraska. Mli: An api for distributed machine learning. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pages 1187–1192. IEEE, 2013.

[18] Qing Wang. A conceptual modeling framework for network analytics. *Data & Knowledge Engineering*, 99:59 – 71, 2015. Selected Papers from the 33rd International Conference on Conceptual Modeling (ER 2014).